Multiscale Applications on European *e*-Infrastructures

Alfons G. Hoekstra & Dieter Kranzlmueller

University of Amsterdam LMU & LRZ, Munich

The Mapper project receives funding from the EC's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° RI-261507.

Two quotes from yesterday

The needs of users of research infrastructures must be better observed already at the planning stage

Norbert Kroo

User engagement is essential, and this is a two-way process

Neil Geddes

Content

- Multiscale Applications
- (Distributed) Multiscale Computing
- The MAPPER project
- Implications for e-Infrastructures

Non-Standard Requirements for e-Infrastructures

- Ability for advanced co-reservation of resources
- Launch emergency simulations
- Consistent interfaces for federated access
- Access to back end nodes: steering, visualisation
- Light path network connections on demand
- Data integration from multiple sources
- Support for software
- Interoperability
- etc.

- MAPPER has written a deliverable(1) on a policy framework needed to support MAPPER applications, and in a broader sense, to support DMC.

(1) Deliverable D3.1, available upon request

Nature is Multiscale

- Natural processes are multiscale
  - 1 H_2O molecule
  - A large collection of H_2O molecules, forming H-bonds
  - Water, and, in solid form, ice.

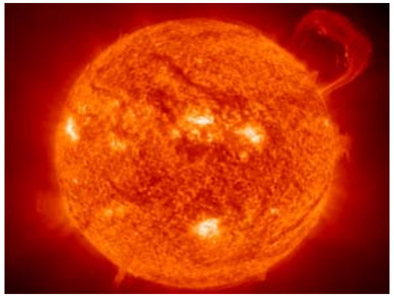

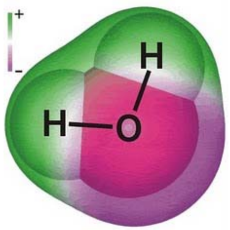

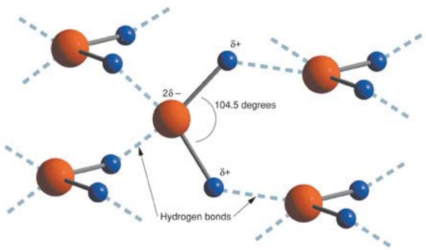

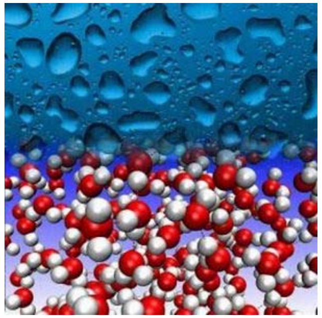

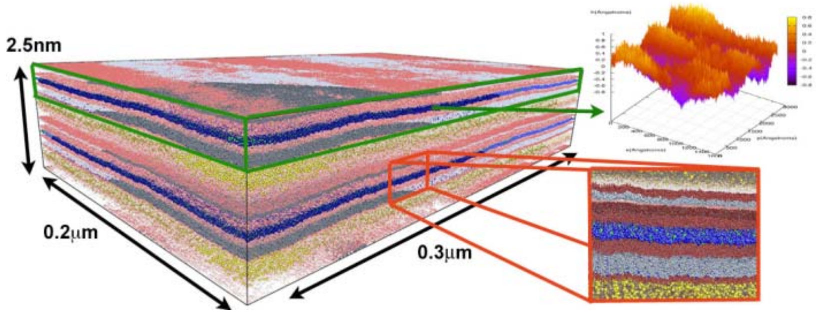

Multiscale models in Biomedicine

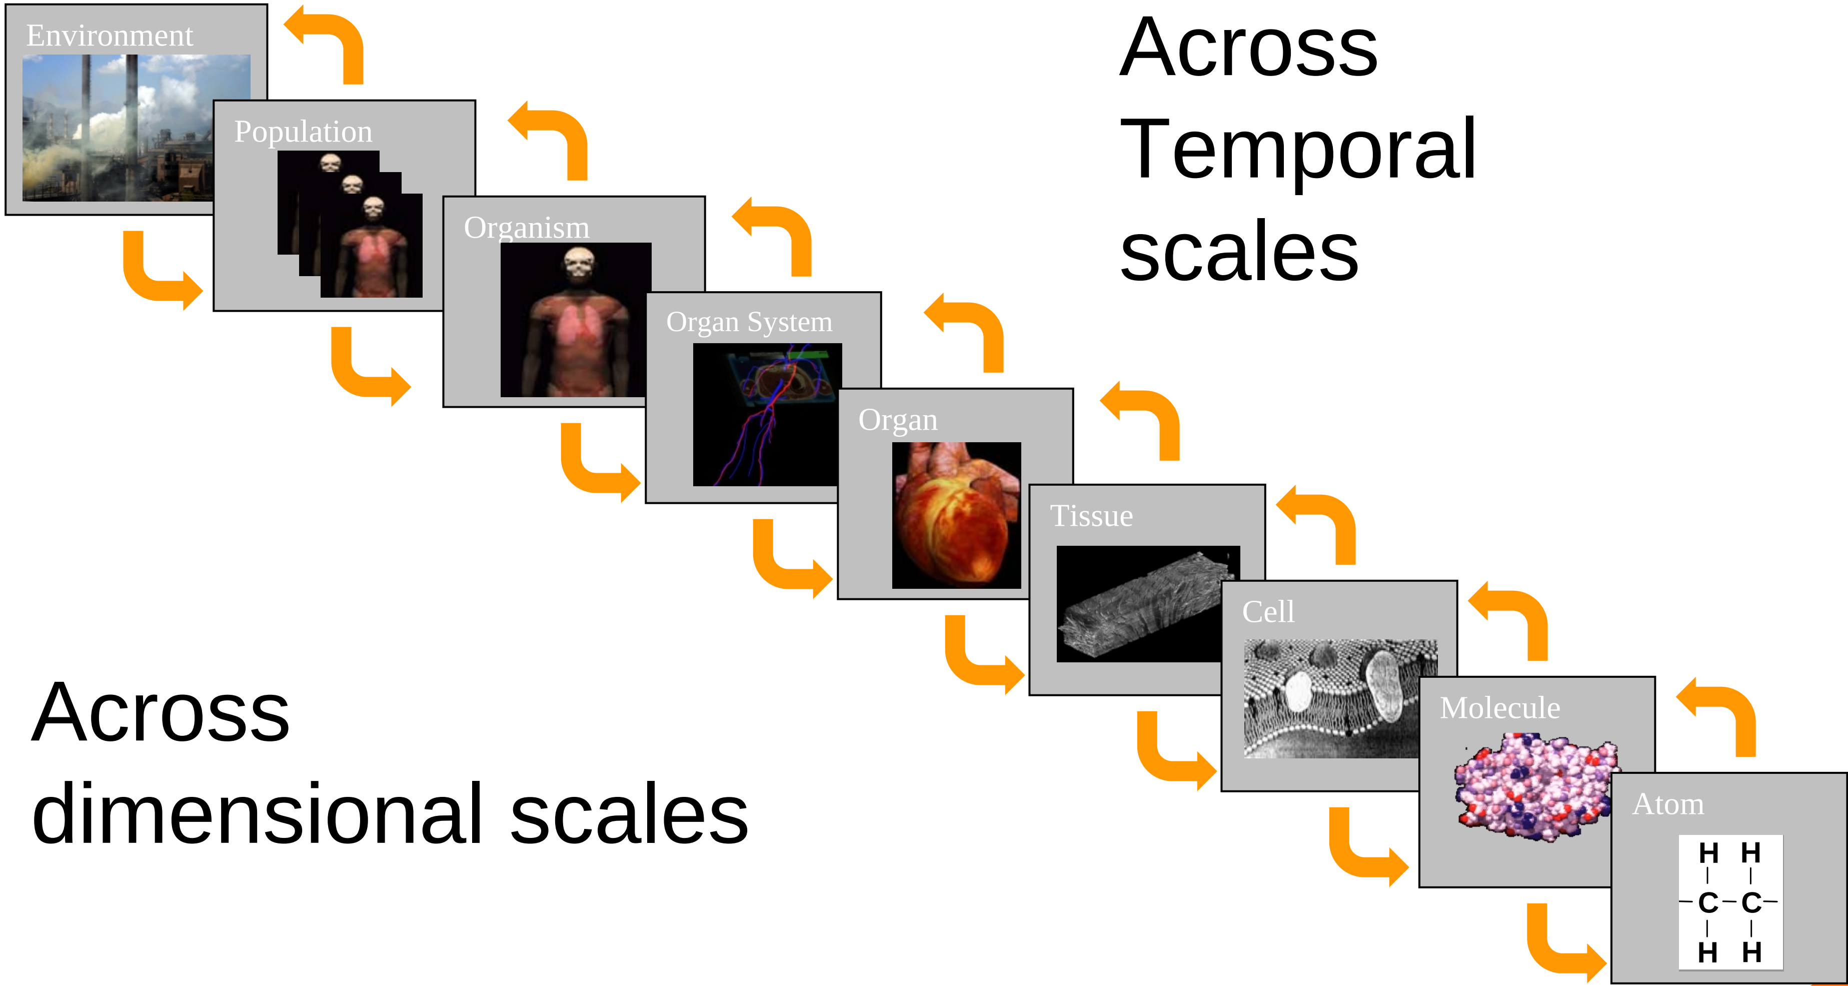

Across Temporal scales

Across dimensional scales

A.G. Hoekstra and P.M.A. Sloot, *Multiscale Biomedical Computing*, Briefings in Bioinformatics 11, 142-152, 2010

Scale range for biomedical applications

- Temporal
  - Molecular events $O(10^{-6})$ s
  - Human life time $O(10^{9})$ s
  - A range of $\mathbf{10^{15}}$
- Spatial
  - Macro molecules $O(10^{-9})$ m
  - Size of human $O(10^{0})$ m
  - A range of $\mathbf{10^{9}}$

Multi-Scale modeling

- Scale Separation Map
- Nature acts on all the scales
- We set the scales
- And then decompose the multiscale system in single scale sub-systems
- And their mutual coupling

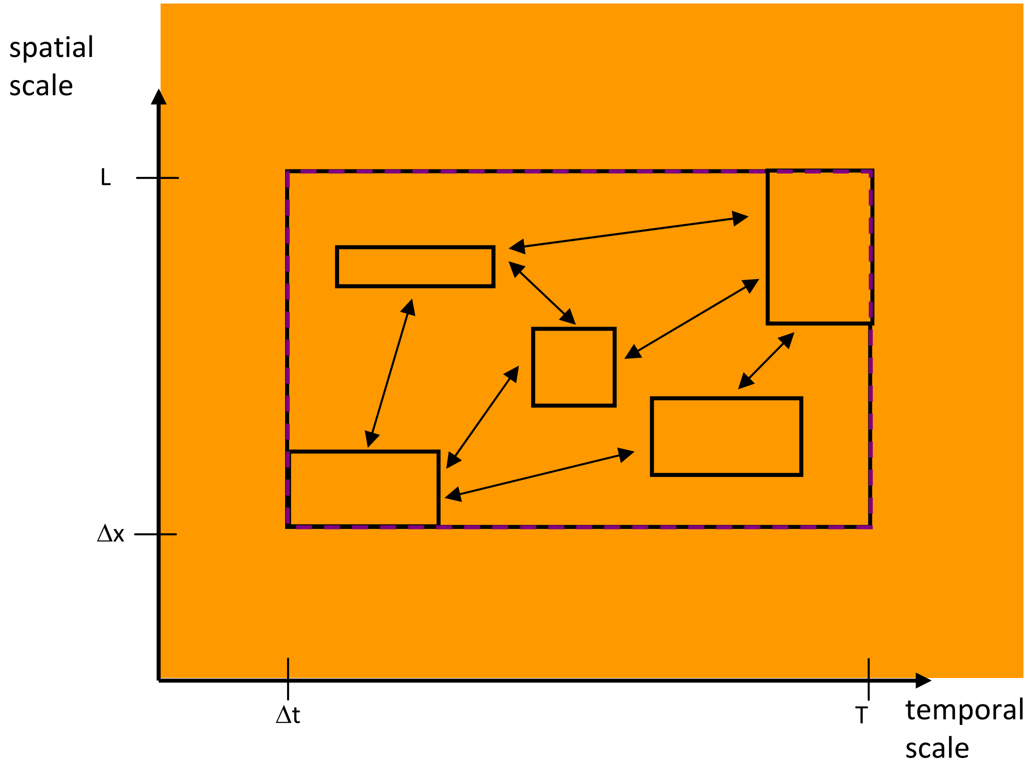

From a Multi-Scale System to many Single-Scale Systems

- Identify the relevant scales
- Design specific models which solve each scale
- Couple the subsystems using a coupling method

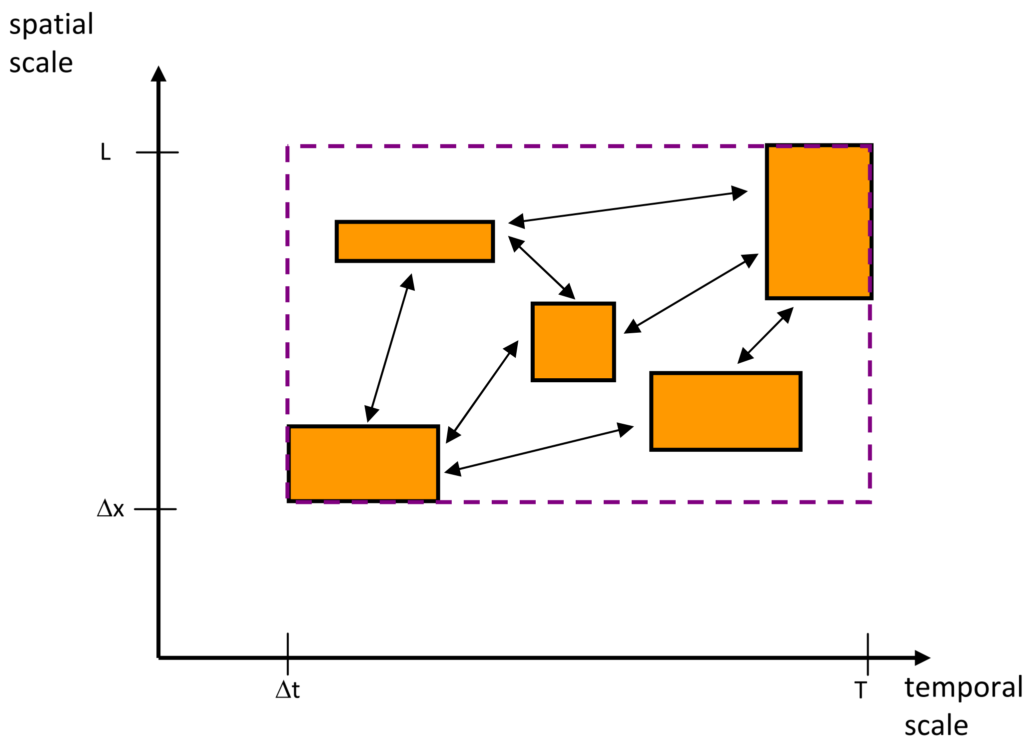

Why multiscale models?

- There is simply no hope to computationally track complex natural processes at their finest spatio-temporal scales.
  - Even with the ongoing growth in computational power.

Minimal demand for multiscale methods

$$\frac{\text{cost of multiscale solver}}{\text{cost of fine scale solver}} << 1$$

$$\text{errors in quantities of interest} < tol$$

Multiscale Speedup

- 1 microscale and one macroscale process
  - At each iteration of the macroscale, the microscale is called
- Execution time full fine scale solver

$$T_{ex}^{full} = \left(\frac{L_M}{\Delta x_\mu}\right)^D \left(\frac{T_M}{\Delta t_\mu}\right)$$

- Execution time for multiscale solver

$$T_{ex}^{multiscale} = \left(\frac{L_M}{\Delta x_M}\right)^D \left(\frac{T_M}{\Delta t_M}\right) \left(\frac{L_\mu}{\Delta x_\mu}\right)^D \left(\frac{T_\mu}{\Delta t_\mu}\right)$$

- Multiscale speedup $S^{multiscale} = \frac{T_{ex}^{full}}{T_{ex}^{multiscale}} = \left(\frac{\Delta x_M}{L_\mu}\right)^D \left(\frac{\Delta t_M}{T_\mu}\right)$

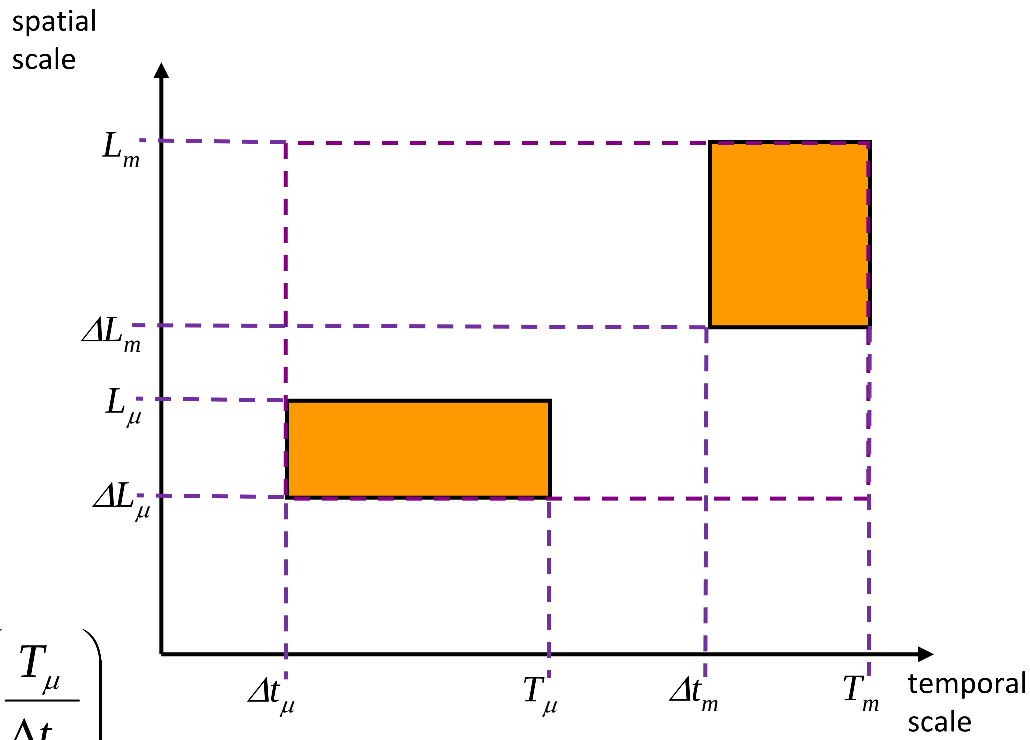

But what about multiscale computing?

- Inherently hybrid models are best serviced by different types of computing environments
- When simulated in three dimensions, they usually require large scale computing capabilities.
- Such large scale hybrid models require a distributed computing ecosystem, where parts of the multiscale model are executed on the most appropriate computing resource.
- Distributed Multiscale Computing

Two Multiscale Computing paradigms

- Loosely Coupled
  - One single scale model provides input to another
  - Single scale models are executed once
  - workflows
- Tightly Coupled
  - Single scale models call each other in an iterative loop
  - Single scale models may execute many times
  - Dedicated coupling libraries are needed

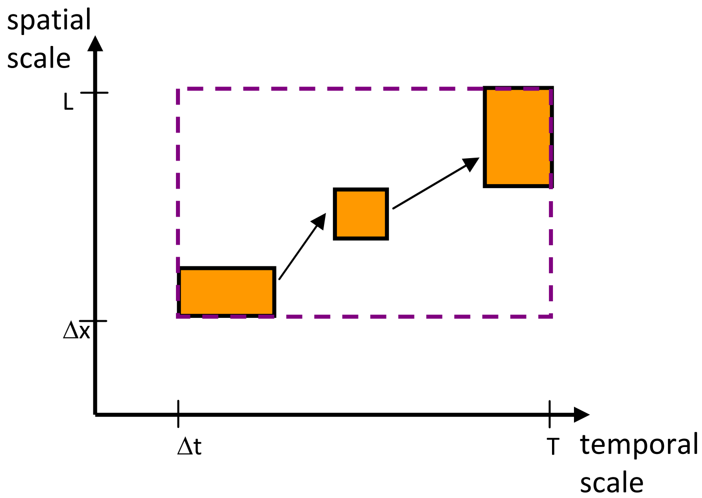

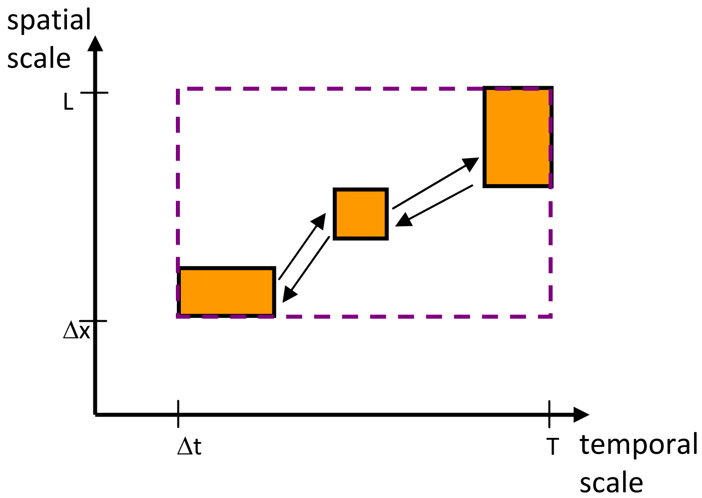

Example: In-stent Restenosis

- Maladaptive response after balloon angioplasty and stenting

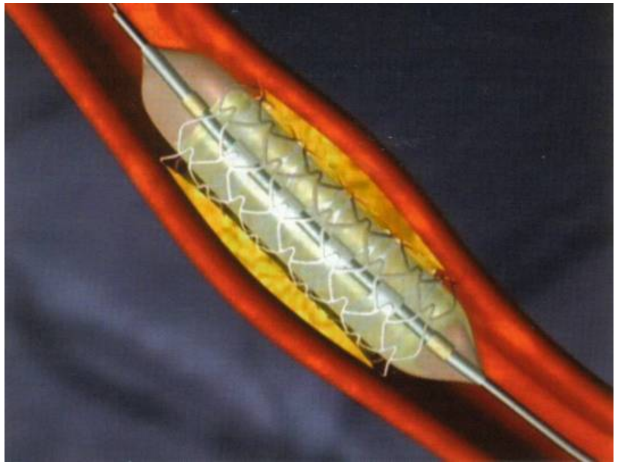

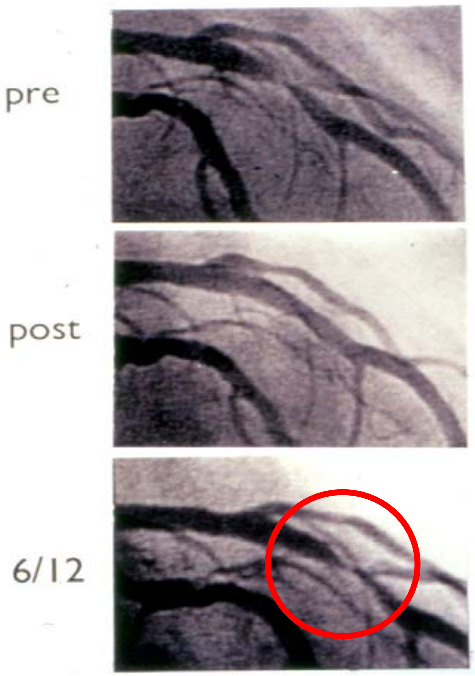

Human angiogram depicting restenosis six months post-PCI.

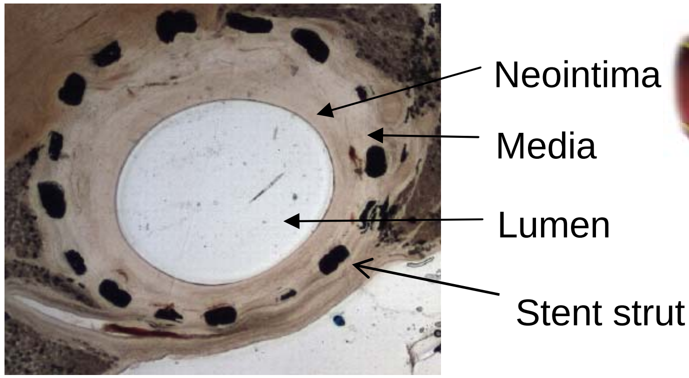

Porcine coronary artery section 28 days post stenting displaying substantial neointima.

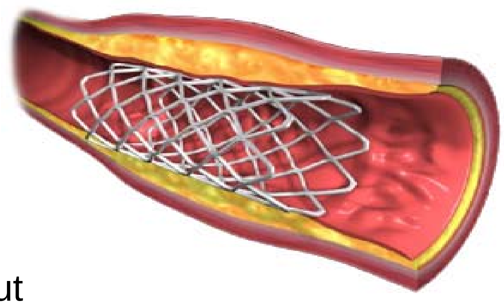

Simplified Scale Separation Map for ISR

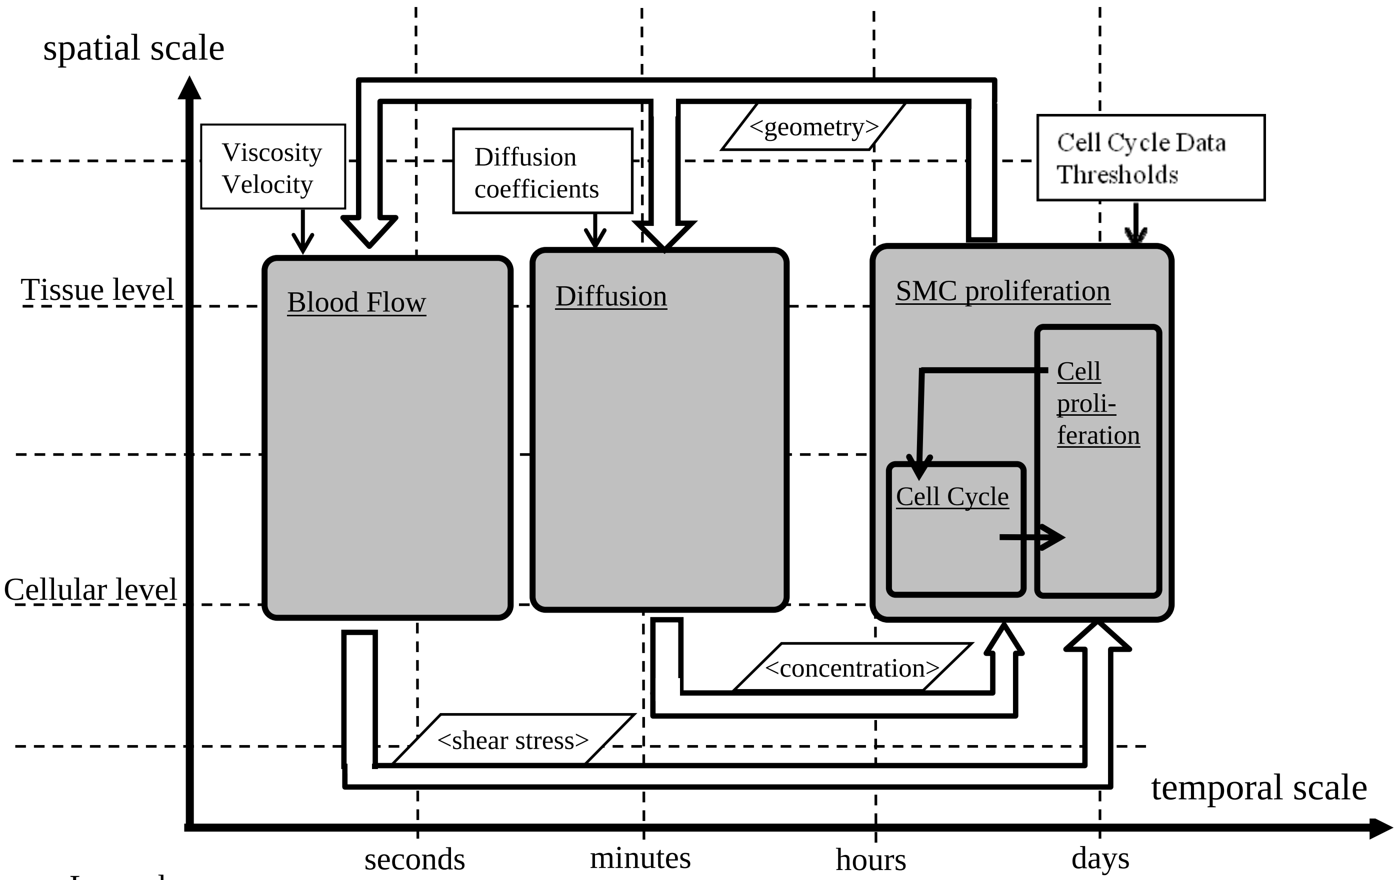

Legend:

Inputs/outputs to single-scale models

Coupling between different-scale models

< … > Data items passed in coupling templates

Some 3D results

SMCs
Stent
Thrombus

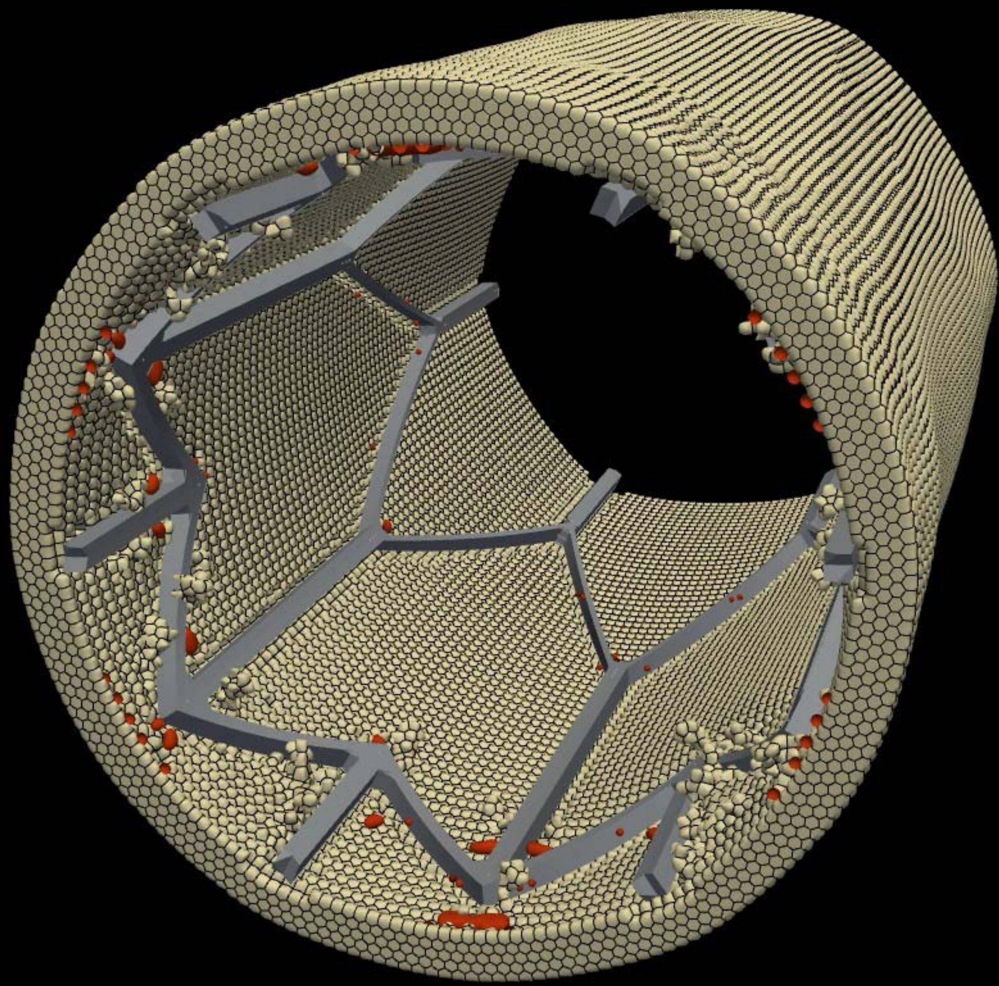

Visualisations:
-- SMC Voronoi tesselation
- fill space with virtual cells
- selective edge smoothing
– Stent: hull triangulation
– Thrombus: isosurfaces

Some 3D results

Drug concentration coloring

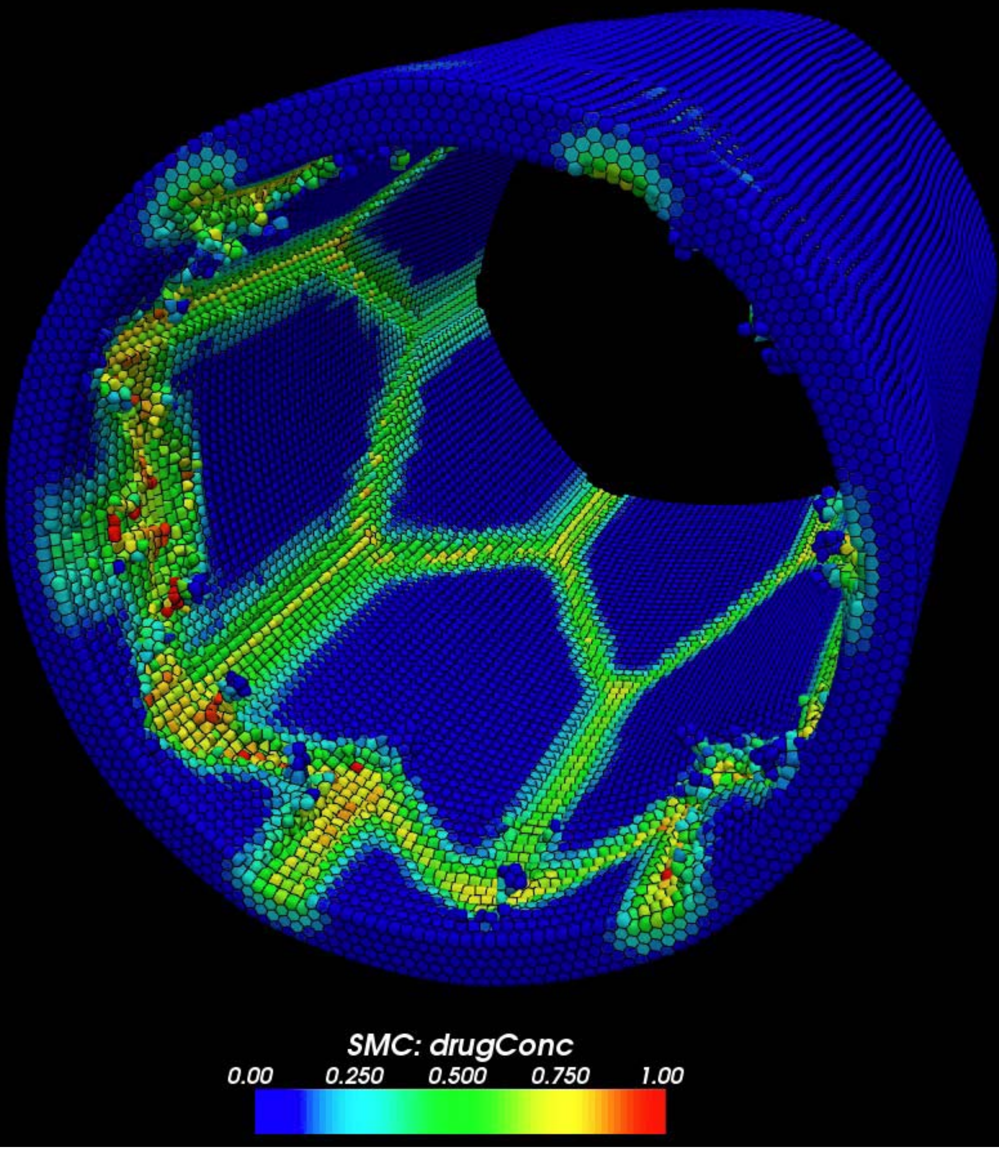

Some 3D results

SMCs (WSS color scale)
Stent
Flow (Ribbons, color scale)

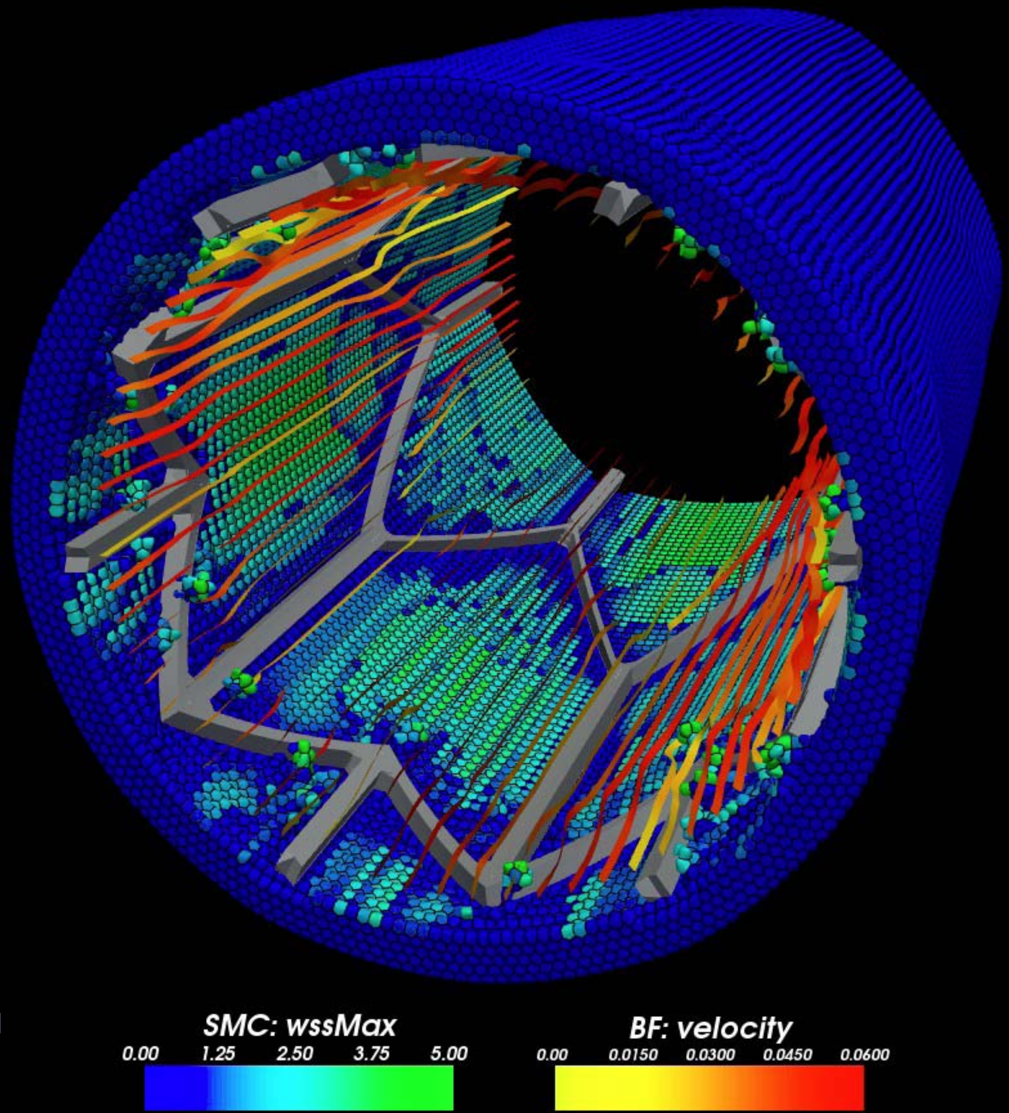

Multiscale APPlications on European e-infRastructures

(proposal number 261507)

www.mapper-project.eu

University of Amsterdam

University College London

University of Ulster

Poznan Supercomputing and Networking Centre

Akademia Gorniczo-Hutnicza im. Stanislawa Staszica w Krakowie

Ludwig-Maximilians-Universität München

University of Geneva

Chalmers Tekniska Högskola

Max-Planck Gesellschaft zur Foerderung der Wissenschaften E.V.

Motivation: user needs

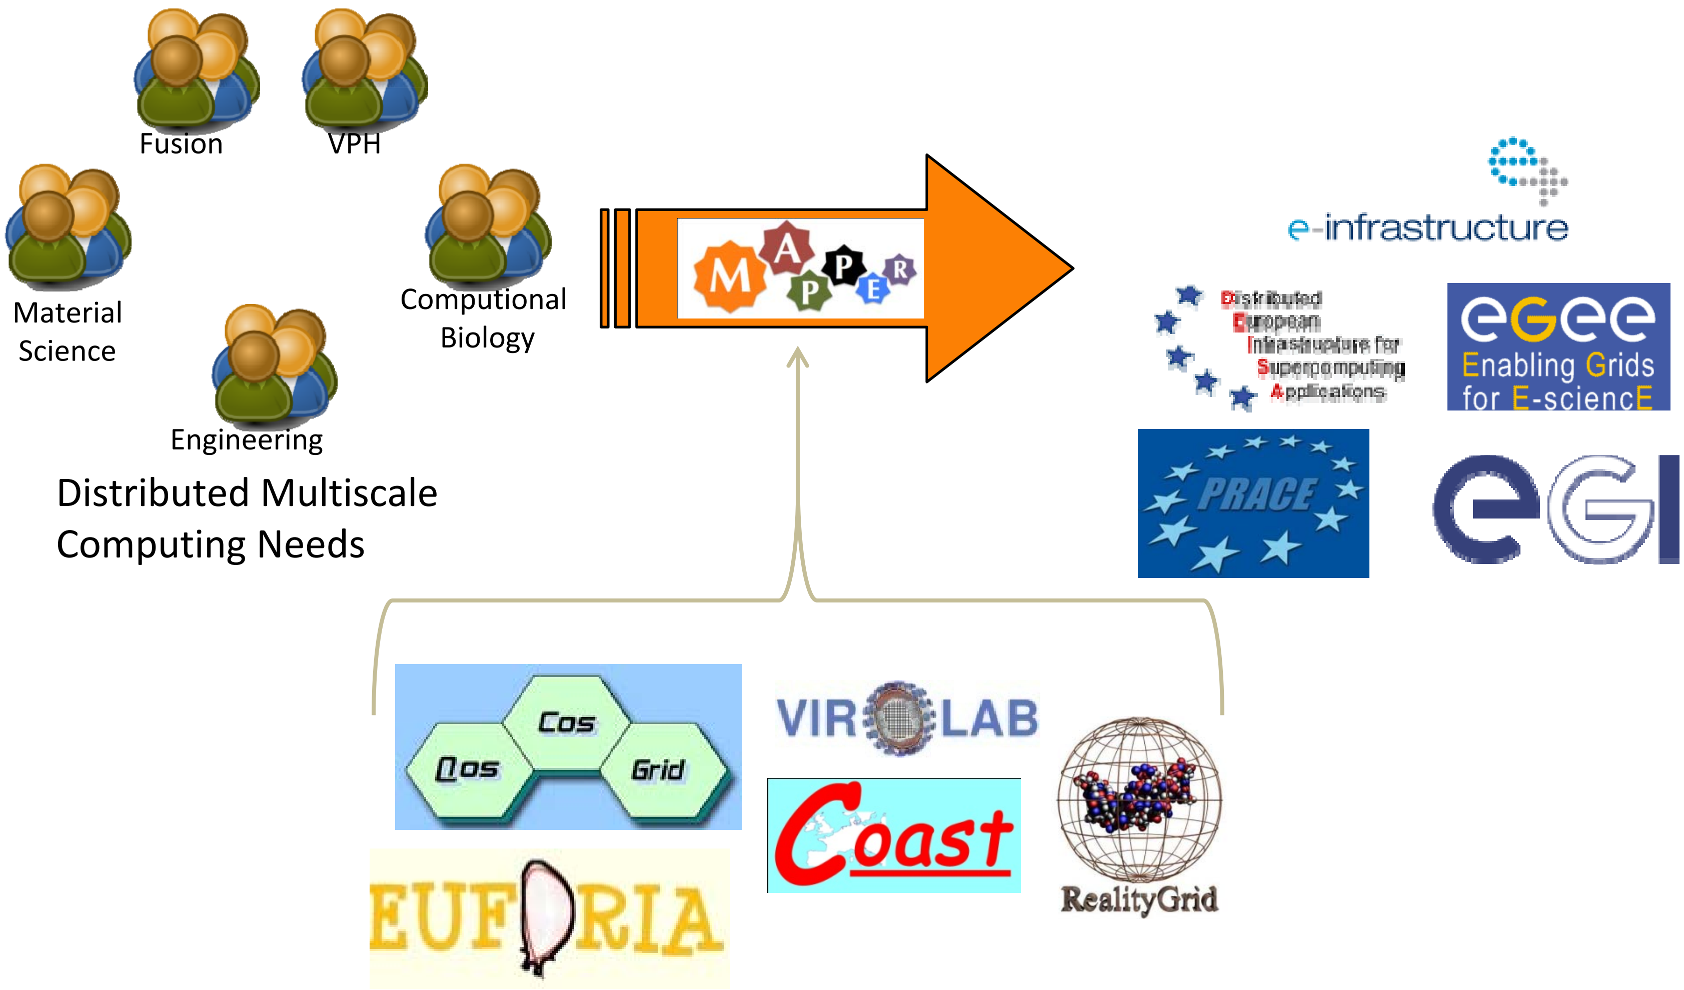

Application Portfolio

virtual physiological human

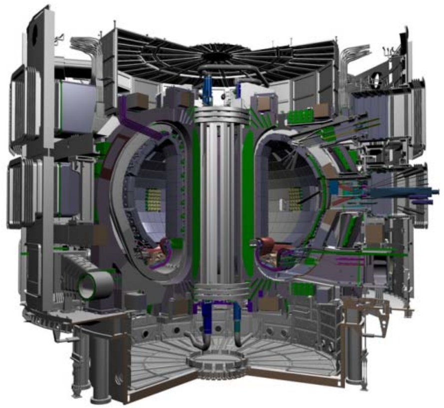

fusion

hydrology

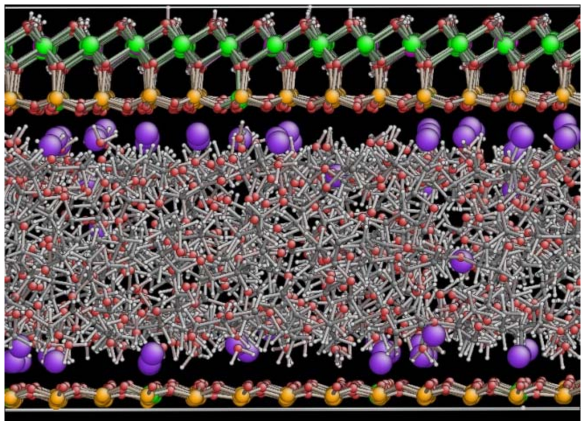

nano material science

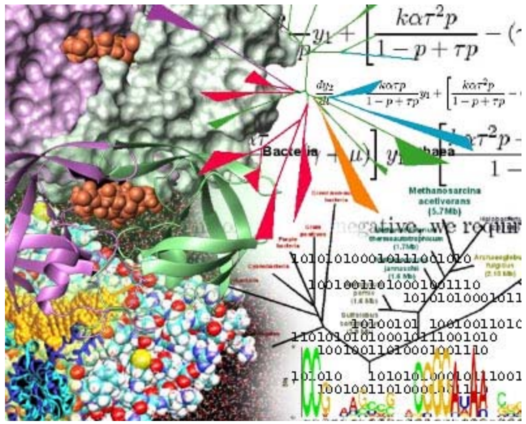

computational biology

Computational power needed

Table 2: Multiscale characteristics of applications

Application	Loosely Coupled	Tightly Coupled	Total number of single scale models	Number of single scale models that require supercomputers
In-stent restenosis		X	5(1)	2
Coupled same-scale and multi-scale hemodynamics		X	3(2)	2
Multi-scale modelling of the BAXS	X		2(3)	1
Edge Plasma Stability	X		3(4)	1
Core Workflow		X	3-10(5)	1-4
Irrigation canals		X	5(6)	1-2
Clay polymers	X		3(7)	2

(1) Blood flow, smooth muscle cell proliferation, drug diffusion, thrombus, stent-deployment; Depending on state-of-the-art when starting the project; (2) HemeLB, a lattice-Boltzmann code for blood flow, NEKTAR, a FEM-based code for blood flow in large arteries, CellML models for cellular processes; (3) metabolism (Phase 1), conjugation (Phase 2) and further modification and excretion (transport) (Phase 3) of the target drug/xenobiotic/endobiotic/bile acid; (4) HELENA or equivalent plasma equilibrium code and ILSA or equivalent plasma stability code; (5) HELENA/CHEASE/EQUAL, some combination of ETAIGB/ NEOWES/ NCLASS/ GLF23/ WEILAND/ GEM, some heating modules from ICRH/NBI/ECRH/LH, some particle source modules from NEUTRALS/PELLETS, some MHD modules from SAWTEETH/NTM/ELMs (6) 1D shallow water models, 2D shallow water models, 2D Free surface flow models, 3D Free surface flow models, Sediment transport models; (7) ab initio molecular dynamics code CASTEP, atomistic molecular dynamics code LAMMPS, coarse-grained simulations also using LAMMPS;

MAPPER Roadmap

- October 1, 2010 – start of project
- Fast track deployment – first year of project
  - Loosely and tightly coupled distributed multiscale simulations can be executed.
- Deep track deployment – second and third year
  - More demanding loosely and tightly coupled distributed multiscale simulation can be executed
  - Programming and access tools available
  - Interoperability available

Again, Implications for e-Infrastructures

- Distributed Multiscale Computing (DMC) leads to non-standard requirements and requests to resource providers.
  - Ability for advance co-reserve of resources
  - Launch emergency simulations
  - Consistent interfaces for federated access
  - Access to back end nodes: steering, visualisation
  - Light path network connections
  - Data integration from multiple sources
  - Support for software
  - Interoperability
  - etc.

- MAPPER has written a deliverable[1] on a policy framework needed to support MAPPER applications, and in a broader sense, to support DMC.

(1) Deliverable D3.1, available upon request

e-infrastructure

MAPPER - 261507 - FP7/2007-2013

CAPACITIES

D3.1 Report on the policy framework resource providers need to adopt to support the MAPPER Project

Project acronym: *MAPPER*

Project full title: Multiscale Applications on European e-Infrastructures

Grant agreement no.: 261507

Due-Date:	M6
Delivery:	M6
Lead Partner:	UCL
Dissemination Level:	CO
Status:	Initial version
Approved:	
Version:	1.0

Available upon request, sent email to

a.g.hoekstra@uva.nl

We face major policy hurdles

- For our projects to be successful, we need integrated compute, storage, networks and services.
  - Current HPC's policies prevent this from happening
    →They still have a batch job mentality!

- No coordinated allocations policies in the EU
  - Need to apply for a project, then if successful apply for compute access
    → Can't do project if compute application rejected!

Importance of connectivity

- With limited national facilities, connectivity to other countries becomes crucial.
- 1-10Gbit wide area networks are needed for large simulations and data movements.
- However, network provisioning is currently extremely difficult and time-consuming.
  - Researchers end up having to request the links, rather than resource providers.

Policy issues

- E-science research has always required changes in resource provider policies to thrive.
- Support for advance machine and network co-reservations.
  - Including urgent computing.
- Improvements in accessibility and usability.
  - Support for Audited Credential Delegation.
  - Interoperability between machines & infrastructures

Political issues

- Streamlined procedures for EU scientific projects.
  - All-in proposals which, when accepted, grant everything needed for a research project.
    - This includes funding for research as well as HPC resource allocations.

- More sensible service level agreements.
  - If a simulation uses multiple machines and one fails, a full allocation refund should be given.

Policy issues 1

- Streamlined procedures for national or EU scientific projects.
  - All-in proposals which, when accepted, grant everything needed for a research project.
    - This includes funding for research as well as HPC resource allocations.

"*To reduce the bureaucratic overhead of EU projects in general, and MAPPER in particular, the procedure of requesting compute time and storage must be greatly streamlined. This can be accomplished by including requests for compute time and storage space in EU project proposals. ….This would eliminate the need for researchers to write multiple proposals for a single project, and prevents projects from receiving a financial budget for research, but not the required computing and storage allocations.*"

- More sensible service level agreements.
  - If a simulation uses multiple machines and one fails, a full allocation refund should be given.

"*When the impact of site and network failures in distributed applications propagate to the full application, and lead to a global breakdown, resource providers should not only refund the hours spent on the crashing site by the application, but the hours spent on the other sites directly involved in execution as well.*"

Policy issues 2

- Support for advance machine and network reservations.
  - Including urgent computing.

"The MAPPER project requires policies which make advance reservations on their resources possible. This applies both to compute infrastructures, as well as wide area network connections. … In the short term, we absolutely require policies that support advance reservation in any way, but to structurally facilitate advance reservation and urgent computing we will need to reach a political agreement with resource providers. This agreement then defines the terms and conditions under which the RPs are willing to support advance reservation and urgent computing. The support for these policies can be arranged per site but is preferably arranged globally through international infrastructure organizations (e.g., PRACE and EGI)."

Policy issues 3

- Improvements in accessibility and usability.
  - Support for Audited Credential Delegation.

"By allowing research groups to do their work using a group certificate in conjunction with Audited Credential Delegation, we can move the overhead of managing grid certificates from the user to the local administrator, and remove one of the largest obstacles for grid accessibility. In addition, Audited Credential Delegation can be used to set up Virtual Organisations."

  - Interoperability between machines & infrastructures.

"RPs should strive for either a uniform stack of middleware and low-level software tools or a completely uniform interface to use these. Doing so is required to achieve interoperability between different compute resources and infrastructures. In addition, we require policies which ensure that job submissions originating from outside the local site are possible."

Policy issues 4

- More advanced issues
  - Advanced co-reservation
    - Obtain multiple resource access at predefined moments in time

"*Advance co-reservation tools will require access to the reservation systems of individual sites. To efficiently support these tools, international organizations such as PRACE and EGI will need to adopt policies to ensure a uniform access interface to local reservation systems. This interface can then be used by co-reservation tools such as the QCG broker or HARC. The reservation of network paths will need to be included in this framework as several MAPPER applications will transport large amounts of data between sites. In compensation for this functionality resource providers could adjust their tariffs so that these "advanced" users pay more per unit of computing time.*"

  - Connectivity policies

"*Site-specific: To accommodate the MAPPER project, a computational site should allow some means for workflow agents and other multi-scale management tools that reside off-site to connect to the local simulation (e.g. by allowing simulations to connect to the outside world under certain conditions). Global: In addition a reservation policy for network connections (including the end-point nodes) is required to deliver a consistent quality of service to the application users.*"

Policy issues 4

- More advanced issues
  - Allocation management

"*A uniform interface policy to access allocation monitoring tools will greatly streamline the allocation management for MAPPER users, as we can then use local software clients to access and obtain all relevant allocation information in one step. Aside from the obvious improvement in usability, this will also make the users more directly aware of their available hours and storage on the sites involved in MAPPER, and prevent them from unknowingly exceeding their allocations.*"

  - Urgent Computing

"*Support for urgent computing should be part of the policy framework of international computing infrastructure organizations such as PRACE and EGI, once the technology has been successfully supported by several supercomputing sites.*"

Concluding remarks

- Multiscale applications
  - prominent in all major domains of science
  - Due to availability of data on all scales
- Distributed Multiscale Computing
  - The natural computing paradigm for many multiscale applications.
- MAPPER will enable DMC on European e-Infrastructures
- Policy issues
  - Technical solutions for DMC exist and are exploited by MAPPER
  - We need however to change a range of policies related to access and use of e-Infrastructures to support DMC.

Suggestion for user engagement

- Co-locate this e-IRG workshop with a major 'computing' conference, such as
  - Supercomputing
    - Yearly, November, USA
  - International Conference on Computational Science
    - Yearly, early June, this year in Singapore
  - E-Science conference
    - Yearly, december, this year in Stockholm

5th-8th December 2011, Stockholm, Sweden
7th IEEE International Conference on e-Science

MAPPER

Alfons Hoekstra & Dieter Kranzlmueller

a.g.hoekstra@uva.nl
kranzlmueller@ifi.lmu.de

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

MNM TEAM
MUNICH NETWORK MANAGEMENT TEAM

University of Amsterdam

Ludwig-Maximilians-Universität München